



Original article

Leveraging prior knowledge for protein–protein interaction extraction with memory network

Huiwei Zhou^{1,*}, Zhuang Liu¹, Shixian Ning¹, Yunlong Yang¹,
Chengkun Lang¹, Yingyu Lin² and Kun Ma³

¹School of Computer Science and Technology, Dalian University of Technology, Chuangxinyuan Building, No. 2 Linggong Road, Ganjingzi District, Dalian, Liaoning, China, ²School of Foreign Languages, Dalian University of Technology, Arts Building, No. 2 Linggong Road, Ganjingzi District, Dalian, Liaoning, China and ³School of Life Science and Medicine, Dalian University of Technology, F03 Building, No. 2 Dagong Road, Liaodongwan District, Panjin, Liaoning, China

*Corresponding author: Tel.: +86 411 84708140; Fax: +86 411 84706706; Email: zhouhuiwei@dlut.edu.cn

Citation details: Zhou,H., Liu,Z., Ning,S. *et al.* Leveraging prior knowledge for protein–protein interaction extraction with memory network. *Database* (2018) Vol. 2018: article ID bay071; doi:10.1093/database/bay071

Received 28 February 2018; Revised 19 May 2018; Accepted 14 June 2018

Abstract

Automatically extracting protein–protein interactions (PPIs) from biomedical literature provides additional support for precision medicine efforts. This paper proposes a novel memory network-based model (MNM) for PPI extraction, which leverages prior knowledge about protein–protein pairs with memory networks. The proposed MNM captures important context clues related to knowledge representations learned from knowledge bases. Both entity embeddings and relation embeddings of prior knowledge are effective in improving the PPI extraction model, leading to a new state-of-the-art performance on the BioCreative VI PPI dataset. The paper also shows that multiple computational layers over an external memory are superior to long short-term memory networks with the local memories.

Database URL: <http://www.biocreative.org/tasks/biocreative-vi/track-4/>

Introduction

With the rapid growth of biomedical literature, it is becoming urgent and significant for natural language processing experts to develop entity–relation extraction techniques (1–4). However, few researches have paid attention to extracting protein–protein interaction affected by mutations (PPI_m) (5). The intricate networks of interactions between genes contribute to controlling cellular homeostasis, and therefore contribute to the development of diseases in specific contexts. Understanding how gene mutations and

variations affect the cellular interactions provides vital support for precision medicine efforts.

For this purpose, the BioCreative VI Track 4 (5) proposes a challenging task of applying biomedical text mining methods to automatically extract interaction relations of protein pairs affected by genetic mutations, which aims to support the precision medicine initiative. There are two specific tasks in Track 4: (i) triage task focuses on identifying scientific abstracts that describe protein–protein interaction (PPI) being disrupted or significantly affected by

genetic mutations and (ii) PPI extraction task focuses on extracting the affected protein pairs. This paper focuses on the PPI extraction task.

This paper presents a novel memory network-based model (MNM) for PPI extraction. The proposed model first encodes the triples (*head entity*, *relation* and *tail entity*) in knowledge bases (KBs) into a continuous vector space, in which a knowledge representation is learned for each entity and relation. Then, the learned knowledge representations are introduced into the memory network through attention mechanisms to capture important context clues toward a pair of entities.

Experiments on the BioCreative VI PPI dataset show that MNM could effectively leverage prior knowledge to improve PPI extraction performance. This paper also shows that multiple computational layers over an external long-term memory are crucial to state-of-the-art performance on the PPI extraction task.

Related work

Previous researches on biomedical relation extraction mostly focus on PPIs (1–4), drug–drug interactions (6–9) and chemical–disease relations (10–14). They can be roughly divided into three categories: rule-based methods, feature-based methods and neural network-based methods.

Rule-based methods extract entity relations by adopting heuristically designed criteria (14, 15). Chen *et al.* (15) assume that a given protein pair contained in more than two sentences within a given document participates in a PPI relationship. Their rule-based system achieves the highest 33.94% *F1*-score on BioCreative VI Track 4 PPI extraction task. Rule-based methods are simple and could achieve good performance on the specific dataset. However, it is hard to apply the extracted rules to a new dataset.

Feature-based methods (1–3, 6–8, 10, 11, 15) apply traditional machine learning techniques to learn models with one-hot represented lexical and syntactic features. Chen *et al.* (15) use Support Vector Machine (SVM) to learn the relation classifier with dependency features and context features. Their feature-based classifier gets the second best reported result (33.66% *F1*-score) in BioCreative VI Track 4 PPI extraction task. Feature-based methods need extensive feature engineering, which is time-consuming and labor intensive.

Recently, deep learning techniques have achieved great success in relation extraction tasks (4, 9, 12, 13, 16–20). Without feature engineering efforts, deep neural networks could effectively extract semantic information for relation extraction. Zeng *et al.* (16) first employed Convolutional Neural Network (CNN) (21) to capture the word and position information for relation extraction, and their model achieves a better performance than feature-based methods.

Tran and Kavuluru (17) employ a CNN to extract local semantic features and get 30.11% *F1*-score on BioCreative VI PPI extraction task. Their system achieves a relatively high precision (36.53%), but suffers from the low recall (25.61%). The reason might be that CNN pays more attention to localized patterns and neglects global dependency information.

A number of recent efforts have been made to capture long-term information within sequences by using Recurrent Neural Network (RNN) (22) or Long Short-Term Memory (LSTM) (23) models. Zhou *et al.* (13) and Zheng *et al.* (9) both use LSTM to model long-distance relation patterns to capture the most important semantic information over a sentence.

Nonetheless, the memory in the LSTM-based models is realized through local memory cells which are locked in the network state from the past, and is inherently unstable over long timescales (24). Weston *et al.* (25) attempt to solve this problem by introducing a class of models called memory networks.

A memory network is a recurrent attention model with a global memory component, which allows being read and written multiple times before outputting a symbol (24). Typically, a memory network consists of a memory m and four components I , G , O , R . m is the input feature representation stored in the memory slot. I converts the input into the internal feature representation. G updates old memories with a new input. O produces an output representation based on the new input and the current memory state. R generates a response according to the output representation. Researches on memory networks show that the multiple computational layers over the long-term memory are crucial for good performance on the tasks of question answering (24, 25) and aspect level sentiment classification (26).

Based on the advantages of memory networks, Feng *et al.* (27) develop a novel attention-based memory network model for relation extraction. Their model consists of a word-level memory network which can learn the importance of each context word with regard to the entity pair, and a relation-level memory network which can capture the dependencies between relations.

All the methods mentioned above use texts as resources. Nevertheless, biomedical experts have built many large-scale KBs, which contain structured triples of protein entity pairs and their interaction relations as the form of (*head entity*, *relation*, *tail entity*), [also denoted as (h , r , t)], such as IntAct (<https://www.ebi.ac.uk/intact/>) (28), BioGrid (<https://thebiogrid.org/>) (29). Both two KBs have the same 45 kinds of PPI relations. Some of relations could be affected by gene mutations such as ‘*physical interactions*’ and ‘*biochemical reactions*’, while some other relations such as ‘*protein complexes*’ and ‘*colocalizations*’ are not considered in the PPI

extraction task. These PPI triples provide a wealth of prior knowledge, which are crucial to PPI extraction.

How to effectively encode this prior knowledge with low-dimensional embeddings of entities and relations is an interesting topic. Recently, knowledge representation learning approach has been proposed to deal with this problem, which attempts to embed the entities and relations into a continuous vector space (30–32). TransE (30) is a typical knowledge representation learning method, which regards a relation r as a translation from the head entity h to the tail entity t with the $\mathbf{h} + \mathbf{r} \approx \mathbf{t}$ in the embedding space, if the triple (h, r, t) holds. Although TransE is very simple, it could achieve state-of-the-art performance on modeling KBs (30). This paper proposes a novel MNM for PPI extraction, which employs TransE to learn embeddings of protein entities and relations from KBs. The learned knowledge representations are then introduced to two memory networks in order to capture important context clues toward a pair of entities. We show that knowledge representations significantly contribute to improving the performance, and the memory network could effectively fuse the prior knowledge and the contextual information.

Materials and methods

Our method for the PPI extraction task can be divided into four steps. Firstly, the candidate instances are generated according to the pre-processing method. Then, entity-relation triples in KBs are fed into TransE model to train embeddings of entities and relations. After that, MNM is employed to capture important context clues related to knowledge representations learned from KBs for PPI relation extraction. Finally, we apply the post-processing rules to find additional PPI relations, and merge them with the results from MNM.

Data

The BioCreative VI Track 4 PPI extraction task corpus contains a total of 2097 PubMed abstracts: 597 for the training set and 1500 for the test set. Proteins in the training corpus are annotated in the form of text offset and length, text span and Entrez Gene ID, only if they participate in mutation affecting PPI relations. PPI relations are annotated at the document level as Entrez Gene ID pair. The number of abstracts and PPI relations in the training and test sets are listed in Table 1.

Pre-processing

Entity recognition and normalization

Protein entities in the training and test sets are recognized by GNormPlus toolkits (<https://www.ncbi.nlm.nih.gov/CB>

Table 1. Statistics of the PPI datasets

Dataset	#Abstract	#PPI
Training	597	752
Test	1500	869
Test-G	1500	483

‘#Abstract’ and ‘#PPI’ mean the number of abstracts and protein–protein interaction affected by mutations in datasets, respectively.

Bresearch/Lu/Demo/tmTools/download/GNormPlus/GNormPlusJava.zip) (33), and normalized to Entrez Gene ID. According to Chen *et al.* (15), GNormPlus have a recall of 53.4%, a precision of 40.5% and an $F1$ -score of 46.1% for the protein name normalization task on the training set, which means not all protein mentions are annotated. For the protein mentions not annotated by GNormPlus in the training set, we can simply put them back based on the annotated protein mentions provided by the training set, and then generate the training protein pairs. But for the protein mentions not annotated by GNormPlus in the test set, we have no evidence to get them back. Table 1 lists the statistics of the test set annotated by GNormPlus (‘Test-G’ for short). There are totally 869 relations in the test set, of which only 483 (55.58%) are remained after protein entity recognition and 386 (44.42%) are lost since the entities in these relations cannot be recognized by GNormPlus. The low recall of entity recognition directly leads to the low recall of relation extraction.

Candidate instances generation

Each abstract in the training set has been manually annotated with at least one relevant interacting protein pair, which is listed with the Entrez Gene ID of the two interactors. If two entities have a PPI relation in a given document, we consider all the mentioned pairs of the two interactors in the document as positive instances.

According to the statistical results of positive instances, we use the following rules to extract candidate instances. Protein pairs not meeting these rules will be discarded.

1. The sentence distance between a protein pair in a candidate instance should be <3 .
2. The token distance between a protein pair should be more than 3 and <50 .

The same rules are applied to the Test-G set. After that, we select the words between a protein pair and three expansion words on both sides of the protein pair as the context word sequence with respect to the protein pair. To simplify the interpretation, we consider the mentions of a protein pair as two single words w_{p_1} and w_{p_2} , where the p_1 and p_2 are the positions of the protein pair. For a given text $\{\dots, w_1, w_2, w_3, w_{p_1}, w_{p_1+1} \dots, w_i, \dots, w_{p_2-1}, w_{p_2}, w_{n-2}, w_{n-1}, w_n, \dots\}$, the context word sequence we

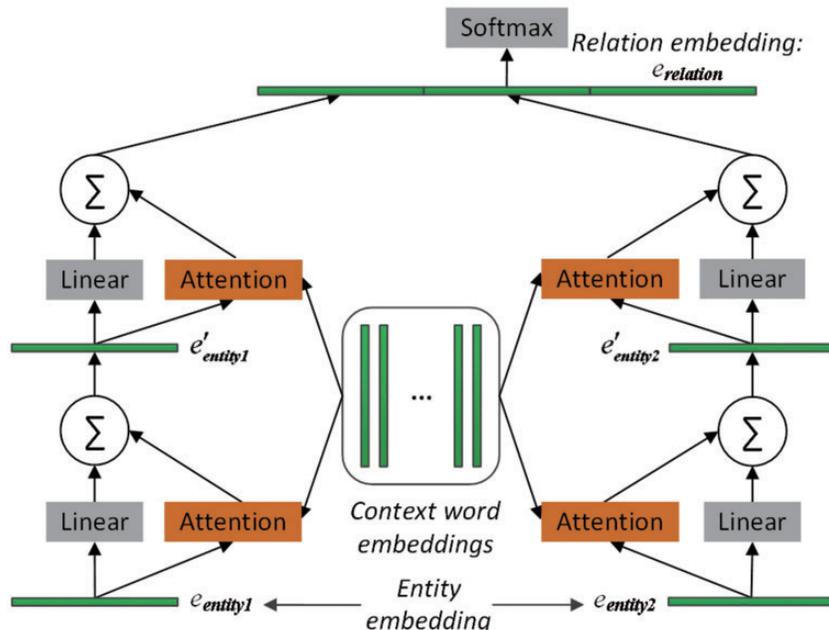


Figure 1. The architecture of the proposed memory network-based model. It consists of two memory networks, each of which contains two computational layers. Embeddings of entity1, entity2 and relation are learned by TransE. Note that the two memory networks share the same parameters, namely, that the same attention operation is applied to both entity1 and entity2. Finally, the two output vectors of the two memory networks and the relation embeddings are concatenated. The resulting vector is fed to the softmax layer for relation classification.

generate is expressed as $\{w_1, w_2, w_3, w_{p_1+1}, \dots, w_i, \dots, w_{p_2-1}, w_{n-2}, w_{n-1}, w_n\}$. As can be seen, we remove the mentions of the protein pair to be classified in the current instance. Then all the other protein mentions are replaced with ‘gene0’. The numbers in the candidate instances are replaced by a specific string ‘NUMBER’. Some special characters, such as ‘*’, are removed.

Finally, the context word sequences of the candidate instances are acquired, which are used as the input of MNM.

Knowledge representation learning

TransE model (30) is employed to learn knowledge representations based on the entity-relation triples in protein KBs IntAct (28) and BioGrid (29). The TransE model regards a relation r as a translation from the head entity h to the tail entity t with the $\mathbf{h} + \mathbf{r} \approx \mathbf{t}$ in the embedding space, if the triple (h, r, t) holds. TransE could learn the structure information from the triples and encode the protein entity embeddings and relation embeddings into a continuous vector space. The loss function of TransE is defined as:

$$L = \sum_{(h,r,t) \in S} \sum_{(h',r,t') \in S'} \max(0, \gamma - \|\mathbf{h} + \mathbf{r} - \mathbf{t}\| + \|\mathbf{h}' + \mathbf{r} - \mathbf{t}'\|) \quad (1)$$

where γ is a margin between correct triples and incorrect triples, S is the set of correct triples and S' is the set of

incorrect triples. KBs only contain correct triples. Conventionally, these correct triples are corrupted by replacing the head or tail entity to generate the incorrect triples (h', r, t) or (h, r, t') . We initialize the entity embeddings with the averaged embeddings of words contained in entity mention, and the relation embeddings with a normal distribution.

Embeddings of entities and relations learned by TransE are introduced into MNM to improve the PPI extraction performance.

Relation extraction

To select the important context words with regard to the pair of proteins, two memory networks are adopted to pay attention to the two entity embeddings, respectively. The architecture of MNM is shown in Figure 1. The two memory networks share the same parameters to learn the weight of each context word of the input sequence. Sharing the same parameters of the attention mechanisms between the two memory networks could enable the two entities to communicate with each other.

In each memory network in Figure 1, there are two computational layers, each of which contains an attention layer and a dimension-wise sum-pooling layer. The outputs of the two networks are concatenated together, and further concatenated with the relation embedding of the protein

pair before given to the softmax layer for relation classification. Next, we will describe MNM in detail.

Attention mechanism

Intuitively, not all the words in the context of a given protein pair describe the PPI relation. For a different entity, the importance of each context word is different as well when we infer the relation of a protein pair. In this work, we employ attention mechanisms to learn the weighted score of each context word with regard to a protein pair. A higher weight indicates higher semantic relatedness with the protein pair.

In each computational layer of the two memory networks, two individual attention mechanisms are adopted to calculate the semantic relatedness of each context word with either of the two entities. The two attention mechanisms share the same parameters. Take one attention mechanism for illustration.

Given a context word sequence $s = \{w_1, w_2, \dots, w_i, \dots, w_n\}$ of a protein pair, the corresponding context word embeddings $\{e_1, e_2, \dots, e_i, \dots, e_n\}$ are regarded as the memory $m \in R^{d \times n}$, where $e_i \in R^d$ is a d -dimensional word embedding, and n is the length of the context word sequence. In each attention layer, each piece of memory m_i is concatenated to one entity embedding to compute its semantic relatedness with the entity. The semantic relatedness score is calculated as follows:

$$g_i = \tanh(W_a[m_i; e_{entity}] + b_a) \quad (2)$$

where $[m_i; e_{entity}]$ denotes the concatenation of memory $m_i \in R^{d \times 1}$ and protein entity embedding $e_{entity} \in R^{d \times 1}$, and $W_a \in R^{1 \times 2d}$ and $b_a \in R^{1 \times 1}$ are attention parameters. After obtaining $\{g_1, g_2, \dots, g_n\}$, the attention weight of each word can be defined as follows:

$$\alpha_i = \frac{\exp(g_i)}{\sum_{j=1}^n \exp(g_j)}. \quad (3)$$

Then the attention layer output $v_{att} \in R^{d \times 1}$ is calculated as a weighted sum of each piece of memory in m :

$$v_{att} = \sum_{i=1}^n \alpha_i m_i. \quad (4)$$

Though the parameters of the two attention mechanisms are the same, the weights of the same context word in the two memory networks are different, since the two concerned entities are different.

By using this attention model, semantic relatedness of each context word with a protein entity can be calculated in an adaptive way. Moreover, this attention model is

differentiable, thus it can be trained easily with other components in an end-to-end fashion.

Dimension-wise sum pooling

The output v_{att} of the attention layer and the linear transformation of entity embedding are fed into a dimension-wise sum pooling layer, and the result vector is considered as a new entity embedding e'_{entity} for the next computational layer:

$$e'_{entity} = W_t e_{entity} \oplus v_{att} \quad (5)$$

where \oplus represents the dimension-wise sum operation, and $W_t \in R^{d \times d}$ is a learned linear transformation matrix. The above sum pooling operation is applied to each of the two entities as shown in Figure 1.

Afterward, the two sum pooling vectors of the last computational layer are concatenated to form the context representations $context = [e'_{entity1}; e'_{entity2}]$. To further take advantage of the prior knowledge, relation embeddings learned from KBs are concatenated to the context representations to form output representations $output_f = [context; e_{relation}]$. Then we pass it to the softmax layer for relation classification.

Position impact

The model mentioned above ignores the position information between context words and entities. Such position information is helpful for attention models because a context word closer to the entities should be more important than a farther one.

Following Sukhbaatar *et al.* (24), we control the input percentage of each piece of memory by its relative distance to the entity mention. Each percentage is calculated as follows:

$$per_i^k = (1 - p_i/n) - (k/d)(1 - 2 \times p_i/n) \quad (6)$$

where n is the input sequence length, k is the number of the current layer, p_i is the relative distance from the current word to the entity mention and d is the dimension of word embeddings. Therefore, the actual memory vector is computed with:

$$m_i^k = e_i \odot per_i^k \quad (7)$$

where \odot represents the dimension-wise product operation.

Relation classification

The softmax layer in Figure 1 consists of a fully connected layer and a logistic regression classifier with a softmax function. It takes $output_f$ as its input and calculates the

probability indicating whether the given protein pair having a PPI relation:

$$p(y = j|T) = \text{softmax}(W_s \text{output}_f + b_s) \quad (8)$$

$$\hat{y} = \underset{y \in \{0,1\}}{\text{argmax}} \left(p(y = j|T) \right) \quad (9)$$

where $W_s \in \mathbb{R}^{2 \times 3d}$ is a learned transformation matrix, b_s is a learned bias vector and T means all training instances.

The cross-entropy loss function is used as the training objective. For each given instance $T^{(l)}$ with its true label $y^{(l)}$, the loss function is calculated as follows:

$$\text{loss} = -\frac{1}{N} \sum_{l=1}^N \log p(y^{(l)}|T^{(l)}) \quad (10)$$

where N is the number of labeled instances in the training set and the superscript l indicates the l -th labeled instance. We adopt Adam technique (34) to update parameters with respect to the loss function.

Post-processing

Inspired by Chen *et al.* (15), we use the following post-processing rules to further improve performance. If more than N sentences refer to a given protein pair within a given document, the protein pair could be considered to have a PPI relation. We set the sentence support threshold N to 2 based on the statistics of the training set.

After extracting positives by the post-processing and MNM, we merge them together as final positives.

In the testing phase, there may exist multiple instances of the same protein pair in a document. It is possible that the same protein pair in a document is predicted inconsistently. If at least one instance is predicted as positive by our model of the same protein pair, we would believe this protein pair has the true PPI relation.

Results and discussion

Experiments are conducted on the BioCreative VI Track 4 PPI extraction task corpus (35). The organizers provide 597 training PubMed abstracts, with the annotated interacting protein pairs (and the corresponding Gene Entrez IDs). The test set consists of 1500 unannotated abstracts which are needed to recognize proteins first and then classify each protein pair into interacting or non-interacting pairs. We directly extract PPI from biomedical documents without document triage (identifying documents that describe PPI impacted by mutations). We train our model by using the training set, and evaluate it on the test set.

Protein entities in the training and test sets are recognized by GNormPlus toolkits, and normalized to Entrez Gene ID. The evaluation of PPI extraction is reported by official evaluation toolkit (<https://github.com/ncbi-nlp/BC6PM>) which adopts micro-averaged (36) Precision (P), Recall (R) and $F1$ -score (F) based on Entrez Gene ID matching (Exact Match) to measure the performance. Note that Micro performance is based on combining results from each interactor protein pair in all documents, weighting equally all pairs, regardless of the number of interactions mentioned in each document.

Word2Vec tool (<https://code.google.com/p/word2vec/>) (37) is used to pre-train word embeddings on the corpus (about 9 308 MB) downloaded from PubMed (<http://www.ncbi.nlm.nih.gov/pubmed/>). The corpus consists of 27 million documents, 3.4 billion tokens and 4.2 million distinct words. TransE model (available at: <https://github.com/thunlp/Fast-TransX>) is employed to learn knowledge representations. The dimension of word, entity and relation embeddings are all 100. If a protein entity is absent in KBs, the entity embedding is initialized as an average of its constituting word embeddings. And for protein pairs not found in KBs, the corresponding relation embeddings are initialized as the zero vector. MNM is trained by using Adam technique (34) with a learning rate 0.001 and a batch size 100. All the hyperparameters are tuned to optimize model by conducting 5-fold cross-validation on the training set. The whole framework is developed by PyTorch (<http://pytorch.org/>).

Prior knowledge resources

We extract PPI relation triples from KBs IntAct (28) and BioGrid (29), and they have the same 45 kinds of PPI. Since protein entities from the two KBs have different identities, we link them to standard database identifiers (Gene Entrez ID in this paper) by using UniProt (38) database. The number of triples extracted from IntAct and BioGrid is shown in Table 2. We merge these two sets of KBs by linking all protein entities to Gene Entrez IDs. Triple selection is strict without redundancy across the two resources. Finally, 1 518 592 unique triples and 84 819 protein

Table 2. Statistics of knowledge bases

Knowledge bases	#Triple	#Protein entity	#Relation type
IntAct	446 992	78 086	45
BioGrid	1 144 450	65 083	45
Merged	1 518 592	84 819	45

The last row named ‘Merged’ means the combination of the triples and protein entities extracted from IntAct and BioGrid. ‘#Triple’, ‘#Protein entity’ and ‘#Relation type’ mean the number of triples, protein entities and relation types, respectively.

entities are obtained for knowledge representation training as shown in Table 2.

The percentage of protein entities and relation triples covered by KBs is shown in Figure 2. From the figure we can see except the relation triple coverage on the test set, KBs cover most of the protein entities and relation triples on the training set and most of the entities on the test set.

From Figure 2, we also find that the training and test sets show different distributions. The main reason is the training set and the test set come from different sources. According to Doğan *et al.* (35), the training set is selected from IntAct (28), which is a specialized curation database including both interaction triples and curated articles. The interaction triples, including PPI, in the IntAct database (28) are derived from curated articles. All these PPI are extracted as prior knowledge resources in this paper. However, the test set contains only novel, not-previously curated articles (35). Therefore, the percentage of entities and triples included in the test set is lower than that in the training set, as shown in Figure 2.

Effects of prior knowledge

In the experiments, we first evaluate the effects of prior knowledge. The proposed MNM with four computational layers is compared with the following baseline methods:

AE (averaged entity embeddings)

This method represents entity embeddings as an average of their constituting word embeddings, and directly feeds the

concatenation of the two networks outputs to the softmax layer. That is to say, entity embeddings and relation embeddings learned from KBs are not used at all.

TE (transE-based entity embeddings)

This method employs TransE-based entity embeddings learned from KBs, and also directly feeds the concatenation of the two networks outputs to the softmax layer. That is to say, only entity embeddings learned from KBs are employed, while relation embeddings are not used.

AE-TR (Averaged Entity Embeddings and TransE-based Relation Embeddings)

This method represents entity embeddings as an average of their constituting word embeddings, and feeds the concatenation of the two network outputs and the relation embeddings to the softmax layer for relation classification. That is to say, entity embeddings learned from KBs are not used, while relation embeddings are employed.

Table 3 lists the comparison results. Seen from the table, MNM outperforms the three baseline methods. Among the three baselines, the best one is AE-TR, which employs relation embeddings learned from KBs. Actually AE-TR is similar to MNM except that the entity embeddings used in AE-TR are not learned from KBs.

Compared with AE, TE employs entity embeddings learned from KBs, and makes the *F1*-score improve by 1.95%, indicating that structured knowledge information contained in TransE-based entity embeddings is more effective than the implicit semantic information expressed by

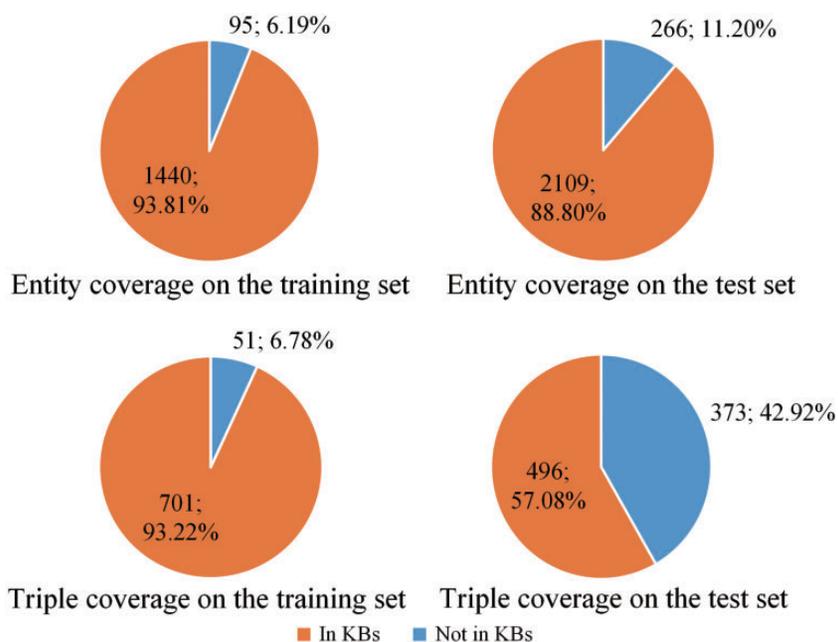


Figure 2. The percentage of protein entities and relation triples covered by knowledge bases. The top two panels show the entity coverage on the training and test sets, respectively. The bottom two panels show the triple coverage on the training and test sets, respectively.

Table 3. Effects of prior knowledge

Prior knowledge	<i>P</i> (%)	<i>R</i> (%)	<i>F</i> (%)
AE	29.86	35.37	32.38**
TE	30.55	39.17	34.33**
AE-TR	37.40	32.14	34.57*
MNM	40.32	32.37	35.91

The bold values are the best results per measure.

AE, TE and AE-TR are the variants of MNM that use prior knowledge or not. The marker * and ** represent *P*-value <0.05 and *P*-value <0.01, respectively, using pairwise *t*-test against MNM.

word embeddings for relation classification. AE-TR simply adds the relation embeddings to AE and improves the *F1*-score by 2.19% compared to AE, which indicates that relation embeddings could provide effective clues about PPI relations to classifier. In MNM, both entity embeddings and relation embeddings are learned from KBs. MNM achieves an *F1*-score of 35.91%, 3.53% higher than AE, and outperforms the top ranked system (15) of this task.

We also show the statistical significance of the overall improvements achieved by our MNM over AE, TE and AE-TR, by using a paired *t*-test. From the results, we can see that all the improvements are statistically significant.

Effects of architecture

To better understand our model, we study three variant architectures of MNM. All the variant architectures have four computational layers and use the same input as MNM.

MNM-Single

This is a single memory network version of MNM, which employs the two entity embeddings in one set of memory units rather than two separate sets of memory units. Specifically, the two entity embeddings are concatenated to each context word embedding to form the input of the attention layer.

MNM-DA

This architecture uses two memory networks as MNM does. However, different from MNM, the attention parameters in the two memory networks are totally different. Here, ‘DA’ is short for ‘different attention’.

MNM-Max

This method does a dimension-wise max pooling operation rather than the dimension-wise sum pooling operation at the end of each layer in MNM.

Table 4 shows the effects of architecture. The observations from Table 4 are listed as follows: (i) the results with two memory networks (MNM-DA, MNM-Max and MNM)

Table 4. Effects of architecture

Architectures	<i>P</i> (%)	<i>R</i> (%)	<i>F</i> (%)
MNM-Single	41.56	30.07	34.89**
MNM-DA	34.70	36.06	35.37*
MNM-Max	38.83	32.83	35.58*
MNM	40.32	32.37	35.91

The bold values are the best results per measure.

MNM-Single, MNM-DA and MNM-Max are the variant versions of MNM. The marker * and ** represent *P*-value <0.05 and *P*-value <0.01, respectively, using paired *t*-test against MNM.

are generally better than the results with a single memory network (MNM-Single); (ii) the two memory networks sharing the same parameters (MNM-Max and MNM) are superior to the two memory networks using the different parameters (MNM-DA). From the results, we could conclude that different attention operations to the two relevant entities would introduce more noises, which are not helpful to relation classification and (iii) comparing with MNM, MNM-Max utilizes dimension-wise max pooling and causes the *F1*-score to drop by 0.33%. It is likely that taking the maximum value of each dimension by max pooling operation may ignore some important contextual information.

Effects of computational layer number

In this section, we further study the effects of the number of computational layers in MNM. Experimental results are listed in Table 5. MNM with different numbers of computational layers are expressed as MNM (*k*), where *k* is the number of the computational layers. When the number is <4, we can observe that more computational layers could generally lead to better performance. The best *F1*-score is achieved when the model contains four computational layers. When computational layer number exceeds four, the performance becomes worse. The reason might lie in the gradient vanishing problem with the number of computational layers increasing.

Comparison with other methods

Table 6 compares our MNM with the following methods:

CNN

This method applies CNN with convolution, max pooling operations. In the convolution layer, 200 feature maps with window size $k = \{3, 4, 5\}$ respectively are learned. The word sequences of CNN are the same as MNM. In order to exploit the position information, this method appends two relative position embeddings to each word embedding in the sequence. And the position is defined as

Table 5. Effects of computational layer number

Computational layer number	P (%)	R (%)	F (%)
MNM(1)	38.26	29.84	33.53
MNM(2)	39.35	31.91	35.24
MNM(3)	33.30	37.21	35.15
MNM(4)	40.32	32.37	35.91
MNM(5)	36.52	34.79	35.63
MNM(6)	34.17	36.18	35.14
MNM(7)	40.95	30.76	35.13
MNM(8)	35.93	35.02	35.47

The bold values are the best results per measure.

The k in MNM(k) is the number of computational layers.

the relative distances from the current word to the head or tail entity.

CNN+KB

In addition to position embeddings, this method appends the two entity embeddings learned from KBs to each context word embedding. Finally, the concatenated representations of the max pooling results and the relation embeddings are fed to the softmax layer.

CNN+KB+Rule

This method merges the protein pairs extracted by CNN+KB and the post-processing rules.

Bi-LSTM

This method applies bi-directional LSTM (23) with both word embeddings and position embeddings. In each direction, the position embedding of the current word toward one of the two entities is concatenated to each word embedding.

Bi-LSTM+KB

In addition to position embeddings, this method appends the two entity embeddings to each context word embedding. For the forward sequence, we concatenate one entity embedding to each word embedding; for the backward sequence, we concatenate the other entity embedding to each word embedding. And in each direction, an attention mechanism is applied to calculate the semantic relatedness of each time step hidden representation with one of the entities. Finally, the bi-directional weighted sum of the hidden representations is concatenated with the relation embeddings. The resulting vectors are fed to the softmax layer for relation classification.

Bi-LSTM+KB+Rule

This method merges the protein pairs extracted by Bi-LSTM+KB and the post-processing rules.

Table 6. Comparison with other methods

Methods	P (%)	R (%)	F (%)
CNN	29.90	34.10	31.86**
CNN+KB	36.02	33.37	34.64*
CNN+KB+Rule	34.75	37.51	36.08*
Bi-LSTM	27.77	38.13	32.14**
Bi-LSTM+KB	38.75	31.57	34.79*
Bi-LSTM+KB+Rule	36.47	36.18	36.32*
MNM	40.32	32.37	35.91
MNM+Rule	37.99	36.98	37.48

The bold values are the best results per measure.

'+KB' means using entity and relation embeddings in the corresponding model. '+Rule' means merging the protein pairs extracted by the post-processing rules and the model. The marker * and ** represent P -value < 0.05 and P -value < 0.01 , respectively, using paired t -test against MNM.

From Table 6, we can see that knowledge representations learned from KBs could consistently improve the performance in both CNN-based and LSTM-based methods, especially the precision of all the methods. As a complement to KBs, the post-processing can improve the recall and achieve a balance between the precision and the recall.

In addition, we find that the LSTM-based models are superior to the CNN-based models in general. This may be due to the fact that the LSTM-based models could capture long-term structure within sequences through local memory cells which are lacking in CNN-based models. After all, PPI relations in this corpus are at document level and mainly reflected in global information.

Furthermore, the proposed MNM-based models perform significantly better than the LSTM-based models. An inherent advantage of MNM-based models lies in the external memory of memory network, which can explicitly reveal the importance of each context word in long sequential data, while LSTM can only capture the contextual information implicitly through local memory cells. Especially for the complex semantic context describing PPI relations, explicitly extracting the important information appears to be more effective.

Comparison with related work

We compare our work with other related work on this task in Table 7. We also report the results evaluated on HomoloGene Match in Table 8. There may be multiple Entrez Gene IDs mapped to the same HomoloGene ID, which causes the difference between Tables 7 and 8. Chen *et al.* (15) apply a rule-based approach which assumes a protein pair contained in more than two sentences within a given document participate in a PPI relationship. This rule-based approach achieves the highest rank in the PPI extraction task. Typically, hand-crafted rules are clear

Table 7. Comparison with related work (Exact Match evaluation)

Related work	P (%)	R (%)	F (%)
CNN (17)	36.53	25.61	30.11
SVM (15)	34.49	32.87	33.66
Rule-based (15)	38.90	30.10	33.94
MNM	40.32	32.37	35.91
MNM+Rule	37.99	36.98	37.48

The bold values are the best results per measure.

'+Rule' means merging the protein pairs extracted by the post-processing rules and the model.

Table 8. Comparison with related work (HomoloGene evaluation)

Related work	P (%)	R (%)	F (%)
CNN (17)	45.44	31.61	37.29
SVM (15)	37.61	35.27	36.40
Rule-based (15)	42.52	33.01	37.17
MNM	42.47	34.22	37.90
MNM+Rule	40.35	39.42	39.88

The bold values are the best results per measure.

'+Rule' means merging the protein pairs extracted by the post-processing rules and the model.

and effective, but they are hard to apply to a new dataset. Chen *et al.* (15) also develop an SVM-based system, which uses dependency features and context features to learn relation classifier. Their feature-based system gets the second best reported result with an *F1*-score of 33.66%. However, feature-based methods need extensive feature engineering, which is time-consuming and labor intensive.

Apart from a traditional rule-based method and machine learning technics, Tran *et al.* (17) employ CNN to implicitly extract semantic features and achieve a relatively high precision. But their approach suffers from the low recall, which is caused by paying attention on localized patterns and neglecting global semantic information. Compared with these systems, our system applies memory networks to fuse contextual information with prior knowledge in KBs. Moreover, our system gets relatively balanced precision and recall after post-processing and outperforms all the systems mentioned above.

Attention visualization

To better demonstrate the effectiveness of attention mechanism, attention weights of two example sequences are visualized in the form of heat maps in Figure 3. In Figure 3, we put the two entity mentions (located on @Entity1 and @Entity2) back to the sequence for clarity, which are removed in practice. For each sequence, the upper and the

lower visible layers show the weights of the context words toward Entity1 and Entity2, respectively. In the first example, 'phosphorylation' and 'kinase' have the maximum weights when we pay attention to the Entity1 and Entity2, respectively. In fact, they are frequent words describing interactions according to Chen *et al.* (15). As for the second example, 'mutant' and 'affinity' have the maximum weights when we pay attention to the Entity1 and Entity2, respectively. According to Chen *et al.* (15), they are usually used to describe gene mutation and interactions, respectively. Figure 4 lists the frequency of top 20 words being assigned the maximum weights in the sequences on the entire test set. From Figure 4, we can observe that key words such as 'phosphorylation', 'mutant', 'bind', 'kinase', 'interact' and 'complex' are in the list. These words are used to indicate mutations or describe interactions frequently according to Chen *et al.* (15). This demonstrates our MNM could identify the important words effectively.

Error analysis

We perform an error analysis the results of MNM+Rule to detect the origins of false positives (FPs) and false negatives (FNs) errors, which are categorized in Figure 5, respectively.

There are two main origins of FPs (shown in the left pie chart of Figure 5):

1. Incorrect classification: 88.75% FPs are from the incorrect classification made by MNM, in spite of the plentiful prior knowledge and detailed contextual information.
2. Rule-based extraction error: post-processing rules cause 11.25% of FPs.

There are three main origins of FNs (shown in the right pie chart of Figure 5):

1. False negative entity: 386 FNs with a proportion of 65.65% are caused by false negative entities without being recognized and normalized by GNormPlus toolkits, which has been mentioned in pre-processing section.
2. Incorrect classification: due to the implicit complex semantic information of protein pairs, MNM misclassifies 147 positive protein pairs as negative.
3. Pre-processing error: protein pairs distributed across more than two sentences are not extracted as candidate instances by pre-processing rules in our system, which causes 55 FNs with a proportion of 9.35%.

Conclusions

This paper develops a novel PPI relation extraction model with two memory networks in order to pay

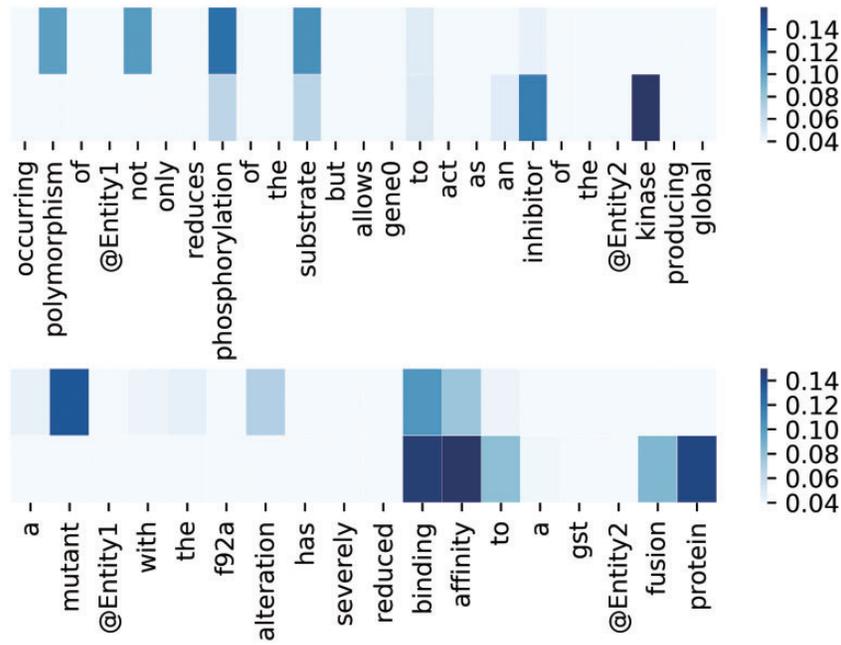


Figure 3. Visualization of attention weights by a heat map. Deeper color means higher weight.

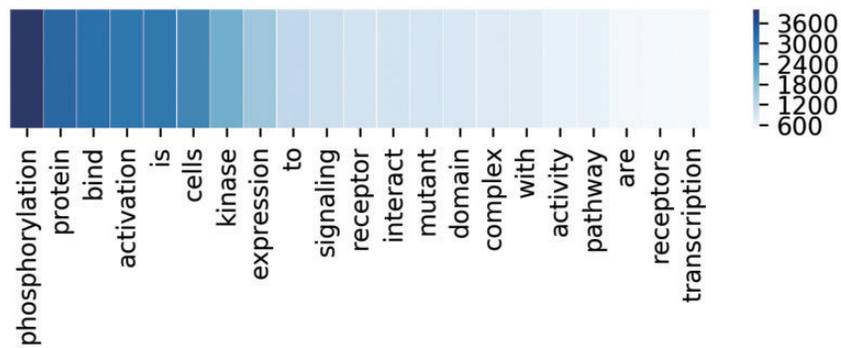


Figure 4. The frequency of top 20 words with the maximum weights in the corresponding sequences. Deeper color means higher frequency.

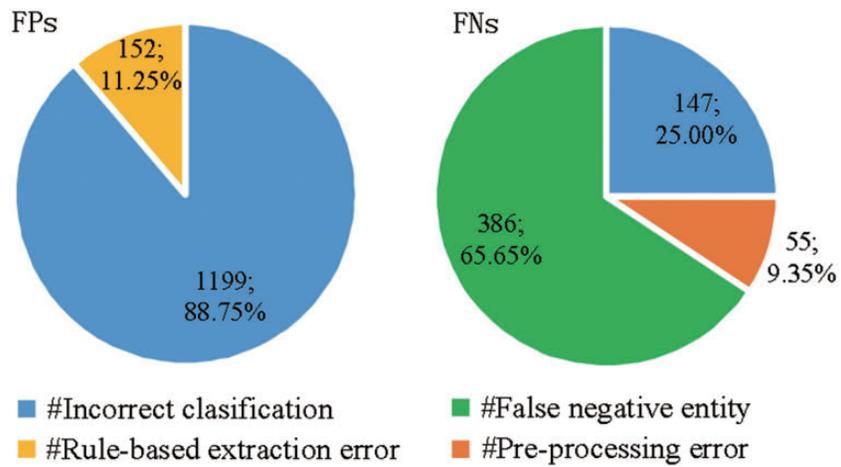


Figure 5. Origins of false positives (FPs) and false negatives (FNs) errors.

attention to the embeddings of protein pairs learned from KBs. The two memory networks share the same parameters and each of memory networks contains multiple computational layers. Experimental result on the BioCreative VI PPI dataset verifies that the proposed approach outperforms the existing state-of-the-art systems. This paper also shows that using multiple computational layers over an external memory is superior to LSTM with local memories.

As future work, we would like to tackle this task at a document level. In this case, how to model the whole document and select the most important information from sentence set in the document is a very challenging problem. We plan to apply a hierarchical attention network to model sentence representations with intra-sentence attention, and then model document representations with inter-sentence attention.

Funding

This work was supported by National Natural Science Foundation of China [grant numbers 61272375, 61772109] and the Ministry of education of Humanities and Social Science research and planning Fund of China [grant numbers 17YJA740076]. Funding to pay the Open Access publication charges for this article was provided by National Natural Science Foundation of China.

Conflict of interest. None declared.

References

- Chowdhury, M.F. and Lavelli, A. (2012) Combining tree structures, flat features and patterns for biomedical relation extraction. In: *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics, Avignon*. pp. 420–429.
- Qian, L. and Zhou, G. (2012) Tree kernel-based protein-protein interaction extraction from biomedical literature. *J. Biomed. Inform.*, **45**, 535–543.
- Phan, T.T. and Ohkawa, T. (2016) Protein-protein interaction extraction with feature selection by evaluating contribution levels of groups consisting of related features. *BMC Bioinformatics*, **17**, 518–246543.
- Peng, Y. and Lu, Z. (2017) Deep learning for extracting protein-protein interactions from biomedical literature. In: *Proceedings of the 2017 Workshop on Biomedical Natural Language Processing, Vancouver*. pp. 29–38.
- Doğan, R.I., Kim, S., Chatr-Aryamontri, A. et al. (2017) Overview of the BioCreative VI precision medicine track. In: *Proceedings of the 2017 Workshop on BioCreative VI, Washington, DC*. pp. 83–87.
- Bui, Q., Sloot, P.M., Van Mulligen, E.M. et al. (2014) A novel feature-based approach to extract drug–drug interactions from biomedical text. *Bioinformatics*, **30**, 3365–3371.
- Kim, S., Liu, H., Yeganova, L. et al. (2015) Extracting drug–drug interactions from literature using a rich feature-based linear kernel approach. *J. Biomed. Inform.*, **55**, 23–30.
- Zhang, Y., Wu, H.Y., Xu, J. et al. (2016) Leveraging syntactic and semantic graph kernels to extract pharmacokinetic drug drug interactions from biomedical literature. *BMC Bioinformatics*, **10**, 67324–67334.
- Zheng, W., Lin, H., Luo, L. et al. (2017) An attention-based effective neural model for drug–drug interactions extraction. *BMC Bioinformatics*, **18**, 445.
- Alam, F., Corazza, A., Lavelli, A. et al. (2016) A knowledge-poor approach to chemical–disease relation extraction. *Database*, **2016**, baw071.
- Pons, E., Becker, B.F., Akhondi, S.A. et al. (2016) Extraction of chemical-induced diseases using prior knowledge and textual information. *Database*, **2016**, baw046.
- Gu, J., Sun, F., Qian, L. et al. (2017) Chemical-induced disease relation extraction via convolutional neural network. *Database*, **2017**, 2.
- Zhou, H., Deng, H., Chen, L. et al. (2016) Exploiting syntactic and semantics information for chemical–disease relation extraction. *Database*, **2016**, baw048.
- Lowe, D.M., Boyle, N.M. and Sayle, R.A. (2016) Efficient chemical–disease identification and relationship extraction using Wikipedia to improve recall. *Database*, **2016**, baw039.
- Chen, Q., Panyam, N.C., Elangovan, A. et al. (2017) Document triage and relation extraction for protein protein interactions affected by mutations. In: *Proceedings of the 2017 Workshop on BioCreative VI, Washington, DC*. pp. 103–106.
- Zeng, D., Liu, K., Lai, S. et al. (2014) Relation classification via convolutional deep neural network. In: *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics, Dublin*. pp. 2335–2344.
- Tran, T. and Kavuluru, R. (2017) Exploring a deep learning pipeline for the BioCreative VI precision medicine task. In: *Proceedings of the 2017 Workshop on BioCreative VI, Washington, DC*. pp. 107–110.
- Zhang, R., Meng, F., Zhou, Y. et al. (2015) Relation classification via recurrent neural network. *arXiv: 1508.01006*, **1**, 234.
- Zhou, P., Shi, W., Tian, J. et al. (2016) Attention-based bidirectional long short-term memory networks for relation classification. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, Berlin*. pp. 207–212.
- Shen, Y. and Huang, X. (2016) Attention-based convolutional neural network for semantic relation extraction. In: *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics, Osaka*. pp. 2526–2536.
- Lecun, Y., Bottou, L., Bengio, Y. et al. (1998) Gradient-based learning applied to document recognition. *Proc. IEEE*, **86**, 2278–2324.
- Schmidhuber, J. (1992) Learning complex, extended sequences using the principle of history compression. *Neural Comput.*, **4**, 234–242.
- Hochreiter, S. and Schmidhuber, J. (1997) Long short-term memory. *Neural Comput.*, **9**, 1735–1780.
- Sukhbaatar, S., Szlam, A., Weston, J. et al. (2015) End-to-end memory networks. In: *Proceedings of the Twenty-Ninth Annual Conference on Neural Information Processing Systems, Montreal*. pp. 2440–2448.
- Weston, J., Chopra, S. and Bordes, A. (2014) Memory networks. *arXiv Preprint arXiv: 1410.3916*.
- Tang, D., Qin, B. and Liu, T. (2016) Aspect level sentiment classification with deep memory network. In: *Proceedings of the 2016*

- Conference on Empirical Methods in Natural Language Processing, *Texas*. pp. 214–224.
27. Feng,X., Guo,J., Qin,B. *et al.* (2017) Effective deep memory networks for distant supervised relation extraction. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne*. pp. 4002–4008.
 28. Hermjakob,H., Montecchialazzi,L., Lewington,C. *et al.* (2004) IntAct: an open source molecular interaction database. *Nucleic Acids Res.*, **32**, 452D–4455.
 29. Stark,C., Breitkreutz,B., Reguly,T. *et al.* (2006) BioGRID: a general repository for interaction datasets. *Nucleic Acids Res.*, **34**, D535–D539.
 30. Bordes,A., Usunier,N., Garciaduran,A. *et al.* (2013) Translating embeddings for modeling multi-relational data. In: *Proceedings of the Twenty-Seventh Annual Conference on Neural Information Processing Systems, Lake Tahoe*. pp. 2787–2795.
 31. Wang,Z., Zhang,J., Feng,J. *et al.* (2014) Knowledge graph embedding by translating on hyperplanes. In: *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, Québec*. pp. 1112–1119.
 32. Lin,Y., Liu,Z., Sun,M. *et al.* (2015). Learning entity and relation embeddings for knowledge graph completion. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, Texas*. pp. 2181–2187.
 33. Wei,C., Kao,H. and Lu,Z. (2015) GNormPlus: an integrative approach for tagging genes, gene families, and protein domains. *BioMed Res. Int.*, **2015**, 1–7.
 34. Kingma,D.P. and Ba,J.L. (2015) Adam: a method for stochastic optimization. In: *Proceedings of the 3rd International Conference for Learning Representation, San Diego*. p. 6.
 35. Doğan,R.I., Chatr-Aryamontri,A., Wei,C.H. *et al.* (2017) The BioCreative VI precision medicine track corpus. In: *Proceedings of the 2017 Workshop on BioCreative VI, Washington*. pp. 88–93.
 36. Krallinger,M., Morgan,A., Smith,L. *et al.* (2008) Overview of the protein-protein interaction annotation extraction task of BioCreative II. *Genome Biol.*, **9**, S1–S19.
 37. Mikolov,T., Sutskever,I., Chen,K. *et al.* (2013) Distributed representations of words and phrases and their compositionality. In: *Proceedings of the Twenty-Seventh Annual Conference on Neural Information Processing Systems, Lake Tahoe*. pp. 3111–3119.
 38. UniProt Consortium. (2015) UniProt: a hub for protein information. *Nucleic Acids Res.*, **43**, D204–D212.